CUR Algorithm with Incomplete Matrix Observation

Rong Jin[†] and **Shenghuo Zhu**[‡]
[†] Dept. of Computer Science and Engineering, Michigan State University, rongjin@msu.edu
[‡] NEC Laboratories America, Inc., zsh@nec-labs.com

arXiv:1403.5647v1 [cs.LG] 22 Mar 2014

1 Introduction

CUR matrix decomposition is a randomized algorithm that can efficiently compute the low rank approximation for a given rectangle matrix (Drineas et al., 2006, Mahoney and Drineas, 2008, 2009). Let $A \in \mathbb{R}^{n\times m}$ be the given matrix and k be the target rank for approximation. CUR randomly samples $c = O(k \log k/\epsilon^2)$ columns and $r = O(k \log k/\epsilon^2)$ rows from A, according to their leverage scores, to form matrices C and R, respectively. The approximated matrix $\widehat{A}$ is then computed as CUR, where $U = C^\dagger A R^\dagger$. It can be shown, that with a high probability,

$$\|A - \widehat{A}\|_F \leq (2+\epsilon)\|A - A_k\|_F \tag{1}$$

where A_k is the best k-rank approximation of A. In case when the maximum of statistical leverage scores, which is also referred to as incoherence measure in matrix completion (Candès and Tao, 2010, Recht, 2011, Candès and Recht, 2012), are small, CUR matrix decomposition can be simplified by uniformly sampling rows an columns from A. The simplified algorithm will have a relative error bound similar to that in (1) except that the sample sizes c and r should be increased by the incoherence measure. In this draft, we will focus on the situation with bounded incoherence measure where uniform sampling of columns and rows is in general sufficient.

One limitation with the existing CUR algorithms is that they require an access to the full matrix A for computing U. In this work, we aim to alleviate this limitation. In particular, we assume that besides having an access to randomly sampled d rows and d columns from A, we only observe a subset of randomly sampled entries Ω from A. Our goal is to develop a low rank approximation algorithm, similar to CUR, based on (i) randomly sampled rows and columns from A, and (ii) randomly sampled entries from A.

Compared to the standard matrix completion theory (Candès and Tao, 2010, Recht, 2011, Candès and Recht, 2012), the key advantage of the proposed algorithm is its low sample complexity and high computational efficiency. In particular, unlike matrix completion that requires $O(rn \log^2 n)$ number of observed entries, the proposed algorithm is able to perfectly recover the target matrix A with only $O(rn \log n)$ number of observed entries (including the randomly sampled entries and entries in randomly sampled rows and columns). In addition, instead of having to solve an optimization problem involved trace norm regularization, the proposed algorithm only needs to solve a standard regression problem. Finally, unlike most matrix completion theories that hold only when the target matrix is of low rank, we show a strong guarantee for the proposed algorithm even when the target matrix A is not low rank.

We finally note that a closely related algorithm, titled "Low-rank Matrix and Tensor Completion via Adaptive Sampling", was published recently (Krishnamurthy and Singh, 2013). It is designed to recover a low rank matrix with randomly sampled rows and entries, which is different from the goal of this work (i.e. computing a low rank approximation for a target matrix A).

2 Algorithm and Notation

Let $M \in \mathbb{R}^{n\times m}$ be the target matrix, where $n \geq m$. To approximate M, we first sample uniformly at random d columns and rows from M, denoted by $A = (\mathbf{a}_1, \ldots, \mathbf{a}_d) \in \mathbb{R}^{n\times d}$ and $B = (\mathbf{b}_1, \ldots, \mathbf{b}_d) \in \mathbb{R}^{m\times d}$, respectively, where each $\mathbf{a}_i \in \mathbb{R}^n$ and $\mathbf{b}_j \in \mathbb{R}^m$ is a row and column of M, respectively. Let r be the target rank for approximation, with $r \leq d$. Let $\widehat{U} = (\widehat{\mathbf{u}}_1, \ldots, \widehat{\mathbf{u}}_r) \in \mathbb{R}^{n\times r}$ and $\widehat{V} = (\widehat{\mathbf{v}}_1, \ldots, \widehat{\mathbf{v}}_r) \in \mathbb{R}^{m\times r}$ be the first r eigenvectors of $AA^\top$ and $BB^\top$, respectively. Besides A and B,

we furthermore sample, uniformly at random, entries from matrix M. Let Ω include the indices of randomly sampled entries. Our goal is to approximately recover the matrix M using A, B, and randomly sample entries in Ω. To this end, we will solve the following optimization problem

$$\min_{Z \in \mathbb{R}^{r \times r}} |\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)|_F^2 \tag{2}$$

where $\mathcal{R}_\Omega : \mathbb{R}^{n \times m} \mapsto \mathbb{R}^{n \times m}$ is defined as

$$[\mathcal{R}_\Omega(M)]_{i,j} = \begin{cases} M_{i,j} & (i,j) \in \Omega \\ 0 & \text{o. w.} \end{cases}$$

Let Z_* be an optimal solution to (2). The recovered matrix is given by $\widehat{M} = \widehat{U} Z_* \widehat{V}^\top$.

The following notation will be used throughout the draft. We denote by $\sigma_i, i = 1, \ldots, m$ the singular values of M in ranked in the descending order, and by $\mathbf{u}_i$ and $\mathbf{v}_i$ be the corresponding left and right singular vectors. Define $U = (\mathbf{u}_1, \ldots, \mathbf{u}_m)$ and $V = (\mathbf{v}_1, \ldots, \mathbf{v}_m)$. Given $r \in [m]$, partition the SVD decomposition of M as

$$M = U \Sigma V^\top = \begin{array}{c} \begin{matrix} r & n-r \end{matrix} \\ \begin{pmatrix} U_1 & U_2 \end{pmatrix} \end{array} \begin{pmatrix} \Sigma_1 & \\ & \Sigma_2 \end{pmatrix} \begin{pmatrix} V_1^\top \\ V_2^\top \end{pmatrix} \tag{3}$$

Let $\widetilde{\mathbf{u}}_i, i \in [n]$ be the ith column of $U_1^\top$ and $\widetilde{\mathbf{v}}_i, i \in [m]$ be the ith column of $V_1^\top$. Define the incoherence measure for U_1 and V_1 as

$$\mu(r) = \max\left(\max_{i \in [n]} \frac{n}{r} |\widetilde{\mathbf{u}}_i|^2, \max_{i \in [m]} \frac{m}{r} |\widetilde{\mathbf{v}}_i|^2 \right)$$

Similarly, we define the incoherence measure for matrices $\widehat{U}$ and $\widehat{V}$. Let $\widehat{\mathbf{u}}_i', i \in [n]$ be the ith column of $\widehat{U}^\top$ and $\widehat{\mathbf{v}}_i', i \in [m]$ be the ith column of $\widehat{V}^\top$. Define the incoherence measure for $\widehat{U}$ and $\widehat{V}$ as

$$\widehat{\mu} = \max\left(\max_{i \in [n]} \frac{n}{r} |\widehat{\mathbf{u}}_i'|^2, \max_{i \in [m]} \frac{m}{r} |\widehat{\mathbf{v}}_i'|^2 \right)$$

Define projection operators $P_U = UU^\top$, $P_V = VV^\top$, $P_{\widehat{U}} = \widehat{U}\widehat{U}^\top$, and $P_{\widehat{V}} = \widehat{V}\widehat{V}^\top$. We will use $\|\cdot\|_2$ for spectral norm of matrix, and $\|\cdot\|_F$ for the Frobenius norm of matrix.

3 Supporting Theorems

In this section, we present several theorems that are important to our analysis.

Theorem 1 *((Halko et al., 2011)) Let M be an $n \times m$ matrix with singular value decomposition $M = U\Sigma V^\top$, an a fixed $r > 0$. Choose a test matrix $\Omega \in \mathbb{R}^{m \times d}$ and construct sample matrix $Y = M\Omega$. Partition M as in (3) and define $\Omega_1 = V_1^\top \Omega$ and $\Omega_2 = V_2^\top \Omega$. Assuming Ω_1 has full row rank, the approximation error satisfies*

$$\|M - P_Y(M)\|_2^2 \leq \|\Sigma_2\|_2^2 + \|\Sigma_2 \Omega_2 \Omega_1^\dagger\|_2^2$$

where $P_Y(M)$ project column vectors in M in the subspace spanned by the column vectors in Y.

Theorem 2 *((Tropp, 2011)) Let $\mathcal{X}$ be a finite set of PSD matrices with dimension k, and suppose that*

$$\max_{X \in \mathcal{X}} \lambda_1(X) \leq B.$$

Sample $\{X_1, \ldots, X_\ell\}$ uniformly at random from $\mathcal{X}$ without replacement. Compute

$$\mu_{\max} = \ell \lambda_{\max}(\mathrm{E}[X_1]), \quad \mu_{\min} = \ell \lambda_{\min}(\mathrm{E}[X_1])$$

Then

$$\Pr\left\{ \lambda_{\max}\left(\sum_{i=1}^{\ell} X_i \right) \geq (1+\delta)\mu_{\max} \right\} \leq k \left[\frac{e^\delta}{(1+\delta)^{1+\delta}} \right]^{\mu_{\max}/B}$$

$$\Pr\left\{ \lambda_{\min}\left(\sum_{i=1}^{\ell} X_i \right) \leq (1-\delta)\mu_{\min} \right\} \leq k \left[\frac{e^{-\delta}}{(1-\delta)^{1-\delta}} \right]^{\mu_{\min}/B}$$

Theorem 3 *Let $A = S^\top HS$ and $\tilde{A} = S^\top \tilde{H} S$ be two symmetric matrices of size $n \times n$. Let $\lambda_i, i \in [n]$ and $\tilde{\lambda}_i, i \in [n]$ be the eigenvalues of A and $\tilde{A}$, respectively, ranked in descending order. Let $U_1, \tilde{U}_1 \in \mathbb{R}^{n\times r}$ include the first r eigenvectors of A and $\tilde{A}$, respectively. Let $\|\cdot\|$ be any invariant norm. Define*

$$\begin{aligned}\Delta_\lambda &= \min\left(\sqrt{2}\left(1 - \frac{\lambda_{r+1}}{\lambda_r}\right), \frac{1}{\sqrt{2}}\right)\\ \Delta_H &= \frac{\|H^{-1}\|_2 \|H - \tilde{H}\|}{\sqrt{1 - \|H^{-1}\|_2 \|H - \tilde{H}\|_2}}\end{aligned}$$

If $\Delta_\lambda \geq \Delta_H/2$, we have

$$\|\sin\Theta(U_1, \tilde{U}_1)\| \leq \frac{\Delta_H}{\Delta_\lambda - \Delta_H/2}\left(1 + \frac{\Delta_H \Delta_\lambda}{16}\right)$$

Since the above theorem follows directly from Theorem 4.4 and discussion in Section 5 from (Li, 1999), we skip its proof.

4 Recovering a Low Rank Matrix

In this section, we discuss the recovery result when the rank of M is no more than r. We will first provide the key results for our analysis, and then present detailed proof for the key theorems.

4.1 Main Result

Our analysis is divided into two steps. We will first show that $|M - P_{\widehat{U}} M P_{\widehat{V}}|_2^2$ is small, and then bound the strongly convexity of the objective function in (2). The following theorem shows that the difference between M and $\widehat{M}$, measured in spectral norm, is well bounded if $|M - P_{\widehat{U}} M P_{\widehat{V}}|_2^2$ is small and the objective function in (2) is strongly convex.

Theorem 4 *Assume (i) $\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq \Delta$, and (ii) the strongly convexity of the objective function is no less than $|\Omega|\gamma$. Then*

$$\|M - \widehat{M}\|_2^2 \leq 2\left(\Delta + \frac{\Delta}{\gamma}\right)$$

To utilize Theorem 4, we need to bound Δ and γ, respectively, which are given in the following two theorems.

Theorem 5 *With a probability $1 - 2e^{-t}$, we have,*

$$\Delta := \|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq 4\sigma_{r+1}^2\left(1 + \frac{m+n}{d}\right)$$

if $d \geq 7\mu(r) r(t + \log r)$.

Proof: Our analysis is based on the following theorem.

Theorem 6 *With a probability $1 - 2e^{-t}$, we have,*

$$\|M - M P_{\widehat{V}}\|_2^2 \leq \sigma_{r+1}^2\left(1 + 2\frac{m}{d}\right), \quad |M - P_{\widehat{U}} M|_2 \leq \sigma_{r+1}^2\left(1 + 2\frac{n}{d}\right)$$

provided that $d \geq 7\mu(r) r(t + \log r)$.

Using Theorem 6, we have, if $d \geq 7\mu(r) r(t + \log r)$, with a probability $1 - 2e^{-t}$

$$\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq 2\|M - M P_{\widehat{V}}\|_2^2 + 2\|(M - P_{\widehat{U}} M) P_{\widehat{V}}\|_2^2 \leq 4\sigma_{r+1}^2\left(1 + \frac{n+m}{d}\right)$$

■

Theorem 7 *With a probability $1 - e^{-t}$, we have that the strongly convexity for the objective function in (2) is bounded from below by $|\Omega|/2$, provided that*

$$|\Omega| \geq 7\widehat{\mu}^2 r^2 (t + 2\log r)$$

The following lemma allows us to replace $\widehat{\mu}$ in Theorem 7 with μ.

Theorem 8 *Assume rank$(M) \leq r$. Then, with a probability $1 - 2e^{-t}$, we have $\widehat{\mu} = \mu(r)$, provided $d \geq 7\mu(r)r(t + \log r)$.*

Proof: When rank$(M) \leq r$, according to Theorem 5, with a probability $1 - 2e^{-t}$, we have $M = P_{\widehat{U}} M P_{\widehat{V}}$, provided that $d \geq 7\mu(r)r(t + \log r)$. Hence $P_{U_1} = P_{\widehat{U}}$ and $P_{V_1} = P_{\widehat{V}}$, which directly implies that $\mu = \widehat{\mu}$. ■

The following theorem follows directly follows from Theorem 5, 7, 8, and 4.

Theorem 9 *Assume rank$(M) \leq r$, $d \geq 7\mu(r)r(t + \log r)$, and $|\Omega| \geq 7\mu^2(r)r^2(t + 2\log r)$. Then, with a probability $1 - 3e^{-t}$, we have $M = \widehat{M}$.*

Remark The result from Theorem 9 shows that, with a probability $1 - \delta$, a low rank matrix M can be perfectly recovered from $O((rn + r^2)\log(r/\delta))$ number of observations from matrix M. This result significantly improves the result from (Krishnamurthy and Singh, 2013), where $O(r^2 n \log(1/\delta))$ number of observations are needed for perfect recovery. We should note that unlike (Krishnamurthy and Singh, 2013) where a small incoherence measure is assumed only for column vectors in matrix M, we assume a small incoherence measure for both row and column vectors in M. It is this assumption that allows us to sample both rows and columns of M, leading to the improvement in the sample complexity.

4.2 Detailed Proofs

4.2.1 Proof of Theorem 4

Set $Z = \widehat{U}^\top M \widehat{V}$. Since $\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq \Delta$, we have

$$\|M - \widehat{U} Z \widehat{V}^\top\|_2^2 \leq \Delta,$$

implying that

$$|M_{i,j} - [\widehat{U} Z \widehat{V}^\top]_{i,j}|^2 \leq \Delta, \ \forall i \in [n], j \in [m]$$

Hence, we have

$$|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)|_F^2 \leq |\Omega|\Delta$$

Let Z_* be the optimal solution to (2). Using the strongly convexity of (2), we have

$$\frac{\gamma}{2}|\Omega|\|Z - Z_*\|_F^2 \leq \frac{1}{2}|\Omega|\Delta,$$

i.e. $\|Z - Z_*\|_F^2 \leq \Delta/\gamma$. We thus have

$$\begin{aligned} \|M - \widehat{M}\|_2^2 &\leq 2\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 + 2\|P_{\widehat{U}} M P_{\widehat{V}} - \widehat{U} Z_* \widehat{V}^\top\|_2^2 \\ &\leq 2\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 + 2\|Z - Z_*\|_2^2 \leq 2\left(\Delta + \frac{\Delta}{\gamma}\right) \end{aligned}$$

4.2.2 Proof of Theorem 6

Let $i_1, \ldots, i_d$ are the d selected columns. Define $\Omega = (\mathbf{e}_{i_1}, \ldots, \mathbf{e}_{i_d}) \in R^{m \times d}$, where $\mathbf{e}_i$ is the ith canonical basis. To utilize Theorem 1, we need to bound the minimum eigenvalue of $\Omega_1 \Omega_1^\top$. We have

$$\Omega_1 \Omega_1^\top = V_1^\top \Omega \Omega^\top V_1$$

Let $\widetilde{\mathbf{v}}_i, i \in [d]$ be the ith row vector of V_1. We have

$$\Omega_1 \Omega_1^\top = \sum_{j=1}^{d} \widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top$$

It is straightforward to show that

$$\mathrm{E}\left[\Omega_1 \Omega_1^\top\right] = \frac{d}{m} I_r$$

To bound the minimum eigenvalue of $\Omega_1 \Omega_1^\top$, we will use Theorem 2. To this end, we have

$$B = \max_{1 \leq i \leq m} |\widetilde{\mathbf{v}}_i|^2 \leq \mu(r)\frac{r}{m}$$

Thus, we have

$$\Pr\left\{\lambda_{\min}(\Omega_1\Omega_1^\top) \le (1-\delta)\frac{d}{m}\right\} \le r \cdot \exp\left(-\frac{d}{\mu(r)r}\left[\delta + (1-\delta)\ln(1-\delta)\right]\right)$$

By setting $\delta = 1/2$, we have, with a probability $1 - re^{-d/[7\mu(r)r]}$

$$\lambda_{\min}(\Omega_1\Omega_1^\top) \ge \frac{d}{2m}$$

Under the assumption that

$$\lambda_{\min}(\Omega_1\Omega_1^\top) \ge \frac{d}{2m},$$

using Theorem 1, we have

$$\|A - AP_{\widehat{V}}\|_2^2 \le \sigma_{r+1}^2 + \left\|\Sigma_2\Omega_2\Omega_1^\dagger\right\|_2^2 \le \sigma_{r+1}^2 + \frac{2m}{d}\|\Sigma_2\Omega_2\|_2^2 \le \sigma_{r+1}^2\left(1 + \frac{2m}{d}\|\Omega_2\|_2^2\right)$$

We complete the proof using the fact that $\|\Omega_2\|_2 \le 1$.

4.2.3 Proof of Theorem 7

We rewrite the objective function as

$$\|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U}Z\widehat{V})\|_F^2 = \left\|\mathcal{R}_\Omega(M) - \sum_{i=1}^r\sum_{j=1}^r Z_{s,t}\mathcal{R}_\Omega(\widehat{\mathbf{u}}_i\widehat{\mathbf{v}}_j^\top)\right\|_F^2$$

Define matrix $K = (\mathbf{k}_1, \ldots, \mathbf{k}_{r^2}) \in \mathbb{R}^{nm\times r^2}$, where $\mathbf{k}_{(i,j)} = \mathrm{vec}\left(\mathcal{R}_\Omega(\widehat{\mathbf{u}}_i\widehat{\mathbf{v}}_j^\top)\right)$. Our goal is to bound the minimum eigenvalue of $K^\top K$. To Theorem 2, we bound

$$B = \max_{i,j}|\mathbf{k}_{(i,j)}|^2 \le \frac{\widehat{\mu}^2 r^2}{mn}$$

and

$$\lambda_{\min}\left(\mathrm{E}[K^\top K]\right) = \frac{|\Omega|}{mn}\lambda_{\min}\left(\widehat{U}\otimes\widehat{V}\right) = \frac{|\Omega|}{mn}$$

where $\otimes$ is Kronecker product. Thus, according to Theorem 2, with a probability $1 - e^{-t}$, we have

$$\lambda_{\min}(K^\top K) \ge \frac{|\Omega|}{2mn}$$

provided that

$$|\Omega| \ge 7\widehat{\mu}^2 r^2(t + 2\log r)$$

5 Recovering the Low Rank Approximation of a Full Rank Matrix

In this section, we consider a general case when M is of full rank but with skewed eigenvalue distribution. To capture the skewed eigenvalue distribution, we use the concept of numerical rank $r(M,\lambda)$ with respect to non-negative constant $\lambda > 0$, which is defined as follows (Hansen, 1987)

$$r(M,\lambda) = \sum_{i=1}^m \frac{\sigma_i^2}{\sigma_i^2 + mn\lambda}$$

Define

$$H_A = \lambda I + \frac{1}{mn}MM^\top, \quad \widehat{H}_A = \lambda I + \frac{1}{dn}AA^\top$$

and

$$H_B = \lambda I + \frac{1}{mn}M^\top M, \quad \widehat{H}_B = \lambda I + \frac{1}{dm}BB^\top$$

Next, we generalize the definition of incoherence measure to numerical low rank. Define $S = \Sigma^2 + mn\lambda I$, where $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_m)$, and incoherence measure μ with respect to a non-negative constant $\lambda > 0$ as

$$\mu(\lambda) = \max\left(\max_{1\le i\le n}\frac{n}{r(M,\lambda)}|V_{i,*}\Sigma S^{-1/2}|^2, \max_{1\le i\le n}\frac{m}{r(M,\lambda)}|U_{i,*}\Sigma S^{-1/2}|^2\right)$$

It is easy to verify that $\mu(\lambda) \ge 1$. Note that when the rank of matrix M is r, we have $r(M,0) = r$ and $\mu(0) = \mu(r)$.

In order to utilize the theorems presented in Section 4 to bound $\|M - \widehat{M}\|_2$, the key is to bound $\mu(r)$ and $\widehat{\mu}$ by $\mu(\lambda)$. The following theorem allows us to bound $\mu(r)$ by $\mu(\lambda)$.

Lemma 1 *Assume*

$$\frac{\sigma_r^2}{(\sigma_r^2 + mn\lambda)r(M,\lambda)} \geq \frac{a}{r}$$

for some positive $a > 0$. We have $\mu(\lambda) \geq a\mu(r)$. More specifically, if we choose $\lambda = \sigma_r^2/mn$, we have

$$\mu(r) \leq \frac{2r(M,\lambda)}{r}\mu(\lambda)$$

Using the above lemma, we have a modified version for Theorem 5

Theorem 10 *Set $\lambda = \sigma_r^2/mn$ for a fixed r. With a probability $1 - 2e^{-t}$, we have,*

$$\Delta := \|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq 4\sigma_{r+1}^2 \left(1 + \frac{m+n}{d}\right)$$

if $d \geq 14\mu(\lambda)r(M,\lambda)(t + \log r)$.

We note that Theorem 10 is almost identical to Theorem 5 except that $\mu(r)r$ is replaced with $\mu(\lambda)r(M,\lambda)$.

Next we will bound $\widehat{\mu}(r)$ by $\mu(\lambda)$. To this end, we need the following theorem.

Theorem 11 *With a probability $1 - 4e^{-t}$, for any $k \in [n]$, we have*

$$1 - \delta \leq \lambda_k(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}) \leq 1 + \delta$$
$$1 - \delta \leq \lambda_k(H_B^{-1/2}\widehat{H}_B H_B^{-1/2}) \leq 1 + \delta$$

if

$$d \geq \frac{4}{\delta^2}(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)$$

Theorem 12 *Assume that $d \geq 16(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)$, and $\sigma_r \geq \sqrt{2}\sigma_{r+1}$. Set $\lambda = \sigma_r^2/mn$. With a probability $1 - 4e^{-t}$, we have*

$$\widehat{\mu}(r) \leq \frac{2r(M,\lambda)}{r}\mu(\lambda) + \frac{18n\delta^2}{r}$$

where

$$\delta^2 = \frac{4}{d}(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)$$

Using Theorem 12, we have the following version of Theorem 7

Theorem 13 *Assume $d \geq 16(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)$ and $\sigma_r \geq \sqrt{2}\sigma_{r+1}$. Set $\lambda = \sigma_r^2/mn$. With a probability $1 - 5e^{-t}$, we have that the strongly convexity for the objective function in (2) is bounded from below by $|\Omega|/2$, provided that*

$$|\Omega| \geq 7\left(2\mu(\lambda)r(M,\lambda) + 72\frac{n}{d}(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)\right)^2 (t + 2\log r)$$

Combining the above results, we have the final theorem for the recovering of M when its numerical rank $r(M,\lambda)$ is small.

Theorem 14 *Assume $d \geq 16(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)$ and $\sigma_r \geq \sqrt{2}\sigma_{r+1}$. Set $\lambda = \sigma_r^2/mn$. We have, with a probability $1 - 7e^{-t}$*

$$\|M - \widehat{M}\|_2^2 \leq 24\sigma_{r+1}^2 \left(1 + \frac{(m+n)}{d}\right)$$

if

$$|\Omega| \geq 7\left(2\mu(\lambda)r(M,\lambda) + 72\frac{n}{d}(\mu(\lambda)r(M,\lambda) + 1)(t + \log n)\right)^2 (t + 2\log r)$$

Remark The total number of observed entries are $\tilde{O}(dn + n^2/d^2)$. It is minimized when $d = n^{1/3}$, leading to $\tilde{O}(n^{4/3})$ for the number of observed entries and $\tilde{O}(\sigma_{r+1}n^{1/3})$ for recovery error.

5.1 Detailed Proof

5.1.1 Proof of Theorem 11

It is sufficient to show the result for $H_A^{-1/2}\widehat{H}_A H_A^{-1/2}$. Define

$$\mathcal{X} = \left\{ M_i = H_A^{-1/2}\left(\frac{1}{n}\mathbf{a}_i\mathbf{a}_i^\top + \lambda I\right) H_A^{-1/2}, i = 1, \ldots, m \right\}$$

We have

$$\begin{aligned} M_i &= US^{-1/2}U^\top (mU\Sigma V_{i,*}^\top V_{i,*}\Sigma U^\top + mn\lambda I) US^{-1/2}U^\top \\ &= U\left(mS^{-1/2}\Sigma V_{i,*}^\top V_{i,*}\Sigma S^{-1/2} + mn\lambda S^{-1}\right)U^\top \end{aligned}$$

Using the definition of $\mu(\lambda)$, we have $\lambda_{\max}(M_i) \le \mu r(M,\lambda) + 1$. Since

$$B = d\lambda_{\max}(\mathrm{E}[M_i]) = d$$

we have

$$\Pr\left\{\lambda_{\max}\left(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}\right) \ge 1+\delta\right\} \le n\exp\left(-\frac{d}{\mu(\lambda)r(M,\lambda)+1}\left[(1+\delta)\log(1+\delta)-\delta\right]\right)$$

Using the fact that

$$(1+\delta)\log(1+\delta) \ge \delta + \frac{1}{4}\delta^2, \forall \delta \in [0,1],$$

we have

$$\Pr\left\{\lambda_{\max}\left(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}\right) \ge 1+\delta\right\} \le n\exp\left(-\frac{d\delta^2}{4(\mu r(M,\lambda)+1)}\right)$$

The upper bound is obtained by setting $d = 4(\mu(\lambda)r(M,\lambda)+1)(\log n + t)/\delta^2$. Similarly, for the lower bound, we have

$$\Pr\left\{\lambda_{\min}\left(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}\right) \le 1-\delta\right\} \le n\exp\left(-\frac{n}{\mu(\lambda)r(M,\lambda)+1}\left[(1-\delta)\log(1-\delta)+\delta\right]\right)$$

Using the fact that

$$(1-\delta)\log(1-\delta) \ge -\delta + \frac{\delta^2}{2}$$

We have the lower bound by setting $m = 2(\mu(\lambda)r(M,\lambda)+1)(\log n + t)/\delta^2$.

5.1.2 Proof of Theorem 12

To utilize Theorem 3, we rewrite H_A and $\widehat{H}_A$, defined in Theorem 11, as

$$H_A = H_A^{1/2} I H_A^{1/2}, \quad \widehat{H}_A = H_A^{1/2} D H_A^{1/2}$$

where $D = H_A^{-1/2}\widehat{H}_A H_A^{-1/2}$. According to Theorem 11, with a probability $1-2e^{-t}$, we have $\|D-I\|_2 \le \delta$, provided that

$$d \ge \frac{4}{\delta^2}(\mu(\lambda)r(M,\lambda)+1)(t+\log n)$$

We then compute Δ_λ and Δ_H defined in Theorem 3. Using the fact $d \ge 16(\mu(\lambda)r(M,\lambda)+1)(t+\log n)$ and Theorem 11, we have, with a probability $1-e^{-t}$, $\delta \le 1/2$. Hence

$$\Delta_H \le \frac{\|D-I\|_2}{\sqrt{1-\|D-I\|_2}} = \frac{\delta}{\sqrt{1-\delta}} \le \sqrt{2}\delta$$

Using the assumption that $\sigma_r/\sigma_{r+1} \ge \sqrt{2}$, we have $\delta \le 1/2 \le 1-\sigma_{r+1}^2/\sigma_r^2$ and therefore $\Delta_\lambda = 1/\sqrt{2}$. As a result, according to Theorem 3, we have

$$\|\sin\Theta(U_1, \widehat{U})\|_2 \le 3\sqrt{2}\delta$$

Similarly, we have,

$$\|\sin\Theta(V_1, \widehat{V})\|_2 \le 3\sqrt{2}\delta$$

Thus, with a probability $1-4e^{-t}$, we have

$$\widehat{\mu}(r) \le \mu(\lambda) + \frac{n}{r}\|\sin\Theta(V_1,\widehat{V})\|_2^2 \le \mu(\lambda) + \frac{18n\delta^2}{r}$$

References

E. Candès and B. Recht. Exact matrix completion via convex optimization. *Commun. ACM*, 55(6): 111–119, 2012.

E. Candès and T. Tao. The power of convex relaxation: near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

P. Drineas, R. Kannan, and M.W. Mahoney. Fast Monte Carlo algorithms for matrices III: Computing a compressed approximate matrix decomposition. *SIAM J Comput*, 36:184–206, 2006.

N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, 2011.

P. C. Hansen. Rank-defficient and discrete ill-posed problems: Numerical aspects of linear inversion. *Society for Industrial and Applied Mathematics*, 1987.

Akshay Krishnamurthy and Aarti Singh. Low-rank matrix and tensor completion via adaptive sampling. In *Adavance in Neural Information Processing (NIPS)*, 2013.

R.-C. Li. Relative perturbation theory: (II) eigenspace and singular subspace variations. *SIAM J. Matrix Anal. Appl.*, 20:471–492, 1999.

M. W. Mahoney and P. Drineas. Relative-error CUR matrix decompositions. *SIAM J Matrix Anal Appl*, 30:844–881, 2008.

M. W. Mahoney and P. Drineas. CUR matrix decompositions for improved data analysis. *Proc. Natl. Acad. Sci. USA*, 106:697–702, 2009.

B. Recht. A simpler approach to matrix completion. *JMLR*, 12:3413–3430, 2011.

J. Tropp. Improved analysis of the subsampled randomized hadamard transform. *Adv. Adapt. Data Anal*, 3:115–126, 2011.